arXiv:1511.08854v2 [cs.CC] 1 Dec 2015

Some Bounds on Communication Complexity of Gap Hamming Distance

Alexander Kozachinskiy*

*Lomonosov Moscow State University, kozlach@mail.ru

December 2, 2015

Abstract

In this paper we obtain some bounds on communication complexity of Gap Hamming Distance problem ($\mathsf{GHD}^n_{L,U}$): Alice and Bob are given binary string of length n and they are guaranteed that Hamming distance between their inputs is either $\leq L$ or $\geq U$ for some $L < U$. They have to output 0, if the first inequality holds, and 1, if the second inequality holds.

In this paper we study the communication complexity of $\mathsf{GHD}^n_{L,U}$ for probabilistic protocols with one-sided error and for deterministic protocols. Our first result is a protocol which communicates $O\left(\left(\frac{s}{U}\right)^{\frac{1}{3}} \cdot n \log n\right)$ bits and has one-sided error probability e^{-s} provided $s \geq \frac{(L+\frac{10}{n})^3}{U^2}$.

Our second result is about deterministic communication complexity of $\mathsf{GHD}^n_{0,\,t}$. Surprisingly, it can be computed with logarithmic precision:

$$\mathrm{D}(\mathsf{GHD}^n_{0,\,t}) = n - \log_2 V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right) + O(\log n),$$

where $V_2(n, r)$ denotes the size of Hamming ball of radius r. As an application of this result for every $c < 2$ we prove a $\Omega\left(\frac{n(2-c)^2}{p}\right)$ lower bound on the space complexity of any c-approximate deterministic p-pass streaming algorithm for computing the number of distinct elements in a data stream of length n with tokens drawn from the universe $U = \{1, 2, \ldots, n\}$. Previously that lower bound was known for $c < \frac{3}{2}$ and for $c < 2$ but with larger $|U|$.

1 Introduction

1.1 Gap Hamming Distance Problem

Given two strings $x = x_1 \ldots x_n \in \{0,1\}^n$, $y = y_1 \ldots y_n \in \{0,1\}^n$, Hamming distance between x and y is defined as the number of positions, where x and y differ:

$$H(x, y) = |\{i \in \{1, \ldots, n\} \,|\, x_i \neq y_i\}|.$$

Let $L < U \leq n$ be integer numbers. In this paper we consider the following communication problem $\mathsf{GHD}^n_{L,U}$, called the Gap Hamming Distance problem:

Definition 1. *Let Alice receive an n-bit string x and Bob an n-bits string y such that either $H(x,y) \leq L$, or $H(x,y) \geq U$. They have to output 0, if the first inequality holds, and 1, if the second inequality holds. If the promise is not fulfilled, they may output anything.*

The Gap Hamming Distance problem is motivated by the problem of approximating the number of distinct elements in a data stream (see [7], [2]). There is the following simple and relatively efficient protocol with shared randomness to solve $\mathsf{GHD}^n_{L,U}$. Alice and Bob pick $i \in \{1,\ldots,n\}$ uniformly at random (using shared randomness) and check, whether $x_i = y_i$ or not. They repeat it many times and then perform some kind of a majority vote: if in more than $\frac{L+U}{2n}$ fraction of trials it happened that $x_i = y_i$ they output 0, and they output 1 otherwise. It can be shown, that $O\left(\frac{snU}{(U-L)^2}\right)$ number of times is sufficient to make error probability less than e^{-s}. Hence

$$R_{e^{-s}}\left(\mathsf{GHD}^n_{L,U}\right) = O\left(\frac{snU}{(U-L)^2}\right). \tag{1}$$

Here $R_\varepsilon(f)$ denotes randomized public-coin communication complexity of f with error probability ε.

Previously the Gap Hamming Distance problem was studied in the symmetric case: $L = \frac{n}{2} - \gamma, U = \frac{n}{2} + \gamma$. Let GHD^n_γ stand for $\mathsf{GHD}^n_{L,U}$ for these specific values of L and U. In this notation the bound (1) becomes $O\left(\frac{n^2}{\gamma^2}\right)$ (for a constant error, say, $\frac{1}{3}$). It turns out that this bound is tight:

Theorem 1 ([4])**.** $R_{\frac{1}{3}}(\mathsf{GHD}^n_\gamma) = \Theta\left(\min\left\{\frac{n^2}{\gamma^2}, n\right\}\right)$.

The most difficult case is $\gamma = c\sqrt{n}$, where c is a constant, in which case the lower bound becomes $R_{\frac{1}{3}}(\mathsf{GHD}^n_{c\sqrt{n}}) = \Omega(n)$. There are several proofs of this bound [4], [11], [9]. As noted in [4], for other values of γ the bound can be proved via the following reduction:

$$R_{\frac{1}{3}}(\mathsf{GHD}^{n/k}_{\gamma/k}) \leq R_{\frac{1}{3}}(\mathsf{GHD}^n_\gamma)$$

for $k > 1$. Setting $k = \Theta\left(\frac{\gamma^2}{n}\right)$ in this inequality, we can reduce Theorem 1 to its special case.

To the best of our knowledge, GHD has not been studied for $L + U \neq n$ (except simple inequality (1)). This paper establishes new bounds on communication complexity of GHD with different parameters and in different settings.

1.2 Our Results

In section 3 we provide the following upper bound on randomized communication complexity of GHD with one-sided error. Before claim it, let us fix our

notations. In this paper $R^i_\varepsilon(f)$ stands for the minimal possible depth of the communication protocol with shared randomness, which never errs on inputs from $f^{-1}(i)$ and which errs with probability at most ε on inputs from $f^{-1}(1-i)$ (here f is partial Boolean function).

Theorem 2. *If* $s \geq \frac{\left(L+\frac{10}{n}\right)^3}{U^2}$, *then*

$$R^0_{e^{-s}}(\mathsf{GHD}^n_{L,\,U}) = O\left(\left(\frac{s}{U}\right)^{\frac{1}{3}} \cdot n \log n\right).$$

Let us compare this bound with the upper bound (1) for protocols with two-sided error. For simplicity assume that $L < \frac{U}{2}$.Then (1) becomes

$$R_{e^{-s}}(\mathsf{GHD}^n_{L,\,U}) = O\left(\frac{s}{U} \cdot n\right).$$

Instead of $\frac{s}{U}$, theorem 2 has $\left(\frac{s}{U}\right)^{\frac{1}{3}}$, which is bigger than $\frac{s}{U}$ when $s < U$. As s tends to U, both $\frac{s}{U}$ and $\left(\frac{s}{U}\right)^{\frac{1}{3}}$ tend to 1 and both bounds become trivial.

In section 4 we study the deterministic communication complexity of $\mathsf{GHD}^n_{0,\,t}$. Namely, we prove the following theorem

Theorem 3.

$$\mathrm{D}(\mathsf{GHD}^n_{0,\,t}) = n - \log_2 V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right) + O(\log n)$$

where $V_2(n, r)$ *denotes the size of Hamming ball of radius* r.

We use this result to prove the following lower bound on space complexity of approximating the number of distinct elements in a data stream:

Theorem 4. *Assume that* $1 < c < 2$ *and* A *is a* p*-pass deterministic streaming algorithm for estimating* F_0*, the number of distinct elements in a given data stream of size* $2n$ *with tokens drawn from the universe* $U = \{1, 2, \ldots, 2n\}$*. If* A *outputs a number* E *such that* $F_0 \leq E < cF_0$*, then* A *must use linear space, namely* $\Omega\left(\frac{n(2-c)^2}{p}\right)$.

Previously such a bound was known in the case when the size of the universe is constant-time larger than the size of the data stream. In the case when the size of the universe and the size of the data stream coincide the bound was known only for $c < \frac{3}{2}$.

2 Preliminaries

2.1 Communication Complexity

Let $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$ be a Boolean function and R an arbitrary random variable whose support is $\mathcal{R}$.

Definition 2. *A randomized (public-coin) communication protocol is a rooted binary tree, in which each non-leaf vertex is associated either with Alice or with Bob and each leaf is labeled by* 0 *or* 1*. For each non-leaf vertex* v*, associated with Alice, there is a function* $f_v : \mathcal{X} \times \mathcal{R} \to \{0,1\}$ *and for each non-leaf vertex* u*, associated with Bob, there is a function* $g_u : \mathcal{Y} \times \mathcal{R} \to \{0,1\}$*. For each non-leaf vertex one of its out-going edges is labeled by* 0 *and other one is labeled by* 1*.*

Definition 3. *Communication complexity of a protocol* π*, denoted by* $CC(\pi)$*, is defined as the depth of the corresponding binary tree.*

A computation according to a protocol runs as follows. Alice is given $x \in \mathcal{X}$, Bob is given $y \in \mathcal{Y}$. They start at the root of tree. If they are in a non-leaf vertex v, associated with Alice, Alice sends $f_v(x, R)$ to Bob and they move to the son of v by the edge labeled by $f_v(x, R)$. If they are in a non-leaf vertex, associated with Bob, they act in a similar same way, however this time it is Bob who sends a bit to Alice. When they reach a leaf, they output the bit which labels this leaf.

Definition 4. *We say that a randomized protocol computes* f *with error probability* ε*, if for every pair of inputs* $(x, y) \in \mathcal{X} \times \mathcal{Y}$ *with probability at least* $1 - \varepsilon$ *that protocol outputs* $f(x, y)$*. Randomized communication complexity of* f *is defined as*

$$R_\varepsilon(f) = \min_\pi CC(\pi),$$

where minimum is over all protocols that compute f *with error probability* ε*.*

If for $i \in \{0,1\}$ we require that the protocol never errs on inputs from $f^{-1}(i)$, then the corresponding notion is called "randomized one-sided error communication complexity" and is denoted by $R^i_\varepsilon(f)$.

If f is a partial function, then, in the definition of computation with error we consider only inputs from the domain of f.

The Gap Hamming Distance problem is the problem of computing the following partial function:

$$\mathsf{GHD}^n_{L,U}(x,y) = \begin{cases} 0 & H(x,y) \le L, \\ 1 & H(x,y) \ge U, \\ \mathsf{undefined} & U < d(x,y) < L, \end{cases} \qquad \text{for } x, y \in \{0,1\}^n.$$

A protocol π is called *deterministic*, if π does not use any randomness.

Definition 5. *We say that a deterministic protocol computes* f*, if for every possible value* $i \in \{0,1\}$ *and for every pair of inputs from* $f^{-1}(i)$ *protocol outputs* i*. Deterministic communication complexity of* f *is defined as*

$$\mathrm{D}(f) = \min_\pi CC(\pi),$$

where minimum is over all deterministic protocols that compute f*.*

2.2 Codes

In section 4 we will use the notion of covering codes.

Definition 6. *A set $C \subset \{0,1\}^n$ is called a* covering code *of radius r, if*

$$\forall x \in \{0,1\}^n \qquad \exists y \in C \qquad H(x,y) \leq r.$$

Obviously, the size a covering code of radius r is at least

$$\frac{2^n}{V_2(n,r)}.$$

There are covering codes with the almost optimal size.

Proposition 1 ([5])**.** *There is a covering code in $\{0,1\}^n$ of radius r and size at most $O\left(\frac{n2^n}{V_2(n,r)}\right)$.*

We will also use the fact that Hamming ball is the largest set among all subsets of $\{0,1\}^n$ with the same diameter.

Definition 7. *Diameter of the set $A \subset \{0,1\}^n$ is equal to*

$$\operatorname{diam}(A) = \max_{x,y \in A} H(x,y).$$

Theorem 5 ([5])**.** *If $B \subset \{0,1\}^n$, $\operatorname{diam}(B) \leq 2r$ and $n \geq 2r+1$, then*

$$|B| \leq V_2(n,r).$$

3 Upper Bound on One-Sided Error Communication Complexity of GHD

Consider any $x, y \in \mathbb{R}^t$. The scalar product and length of a vector are defined in the usual way

$$\langle x, y \rangle = \sum_{i=1}^{t} x_i y_i, \qquad \|x\| = \sqrt{\langle x, x \rangle}.$$

Let $\mathcal{US}(t)$ denote the uniform distribution on $(t-1)$–dimensional unit sphere.

Proposition 2 ([8])**.** *$\mathcal{US}(t)$ is equal to the distribution of the following vector*

$$\frac{(Z_1, Z_2, \ldots, Z_t)}{\sqrt{Z_1^2 + \ldots + Z_t^2}},$$

where $Z_1, \ldots Z_t$ are independent random variables and for each of them we have that $Z_i \sim \mathcal{N}(0,1)$.

Lemma 1. *If $Z \sim \mathcal{US}(t)$, then for each $x \in \mathbb{R}^t$ we have*

$$\mathrm{E}\langle x, Z\rangle^2 = \frac{\|x\|^2}{t}.$$

Proof. Let $Z_1, \ldots, Z_t$ be random variables from Proposition 2. Assume that $x = (1, 0, \ldots, 0)$. Then we have

$$\langle x, Z\rangle^2 = \frac{Z_1^2}{Z_1^2 + \ldots + Z_t^2}.$$

Random variables

$$\frac{Z_1^2}{Z_1^2 + \ldots + Z_t^2}, \frac{Z_2^2}{Z_1^2 + \ldots + Z_t^2}, \ldots, \frac{Z_t^2}{Z_1^2 + \ldots + Z_t^2}$$

are identically distributed. Hence

$$1 = \mathrm{E}\frac{Z_1^2 + \ldots + Z_t^2}{Z_1^2 + \ldots + Z_t^2} = t\mathrm{E}\frac{Z_1^2}{Z_1^2 + \ldots + Z_t^2} = t\mathrm{E}(x, Z)^2.$$

Thus lemma is proved for $x = e_1 = (1, 0, \ldots, 0)$.

Consider any other $x \in \mathbb{R}^t$. If $x = 0$, lemma is obvious. Otherwise there exists an orthogonal $n \times n$ matrix A such that $\frac{x}{\|x\|} = Ae_1$. Now consider the vector $A^T Z$. Proposition 2 implies that vectors $A^T Z$ and Z are identically distributed. Hence

$$\mathrm{E}\langle x, Z\rangle^2 = \|x\|^2 \mathrm{E}\langle Ae_1, Z\rangle^2 = \|x\|^2 \mathrm{E}\langle e_1, A^T Z\rangle^2 = \|x\|^2 \mathrm{E}\langle e_1, Z\rangle^2 = \frac{\|x\|^2}{t}.$$

□

Now we are able to construct the protocol for Theorem 2.

Proof of Theorem 2. Set $b = \left\lceil 4n \sqrt[3]{\frac{s}{U}} \right\rceil$. If $b > n$, then the theorem 2 states that $R^0_{e^{-s}}(\mathsf{GHD}^n_{U,L})$ is linear in n, which is trivial. Therefore we will assume that $b \leq n$. Communication complexity of the protocol will be $O(b \log n)$. Set $a = \left\lceil \frac{n}{b} \right\rceil$ and

$$x_{n+1} = \ldots = x_{ab} = y_{n+1} = \ldots = y_{ab} = 0.$$

Note that

$$ab = \left\lceil \frac{n}{b} \right\rceil b \geq \frac{n}{b} \cdot b = n, \qquad ab = \left\lceil \frac{n}{b} \right\rceil b \leq \left(\frac{n}{b} + 1 \right) b = n + b \leq 2n. \quad (2)$$

Alice and Bob transform their inputs x, y to vectors $\alpha, \beta \in \mathbb{R}^{ab}$, where

$$\alpha = (x_1, \ldots, x_n, x_{n+1}, \ldots, x_{ab}), \qquad \beta = (y_1, \ldots, y_n, y_{n+1}, \ldots, y_{ab}).$$

Note that $H(x,y) = \|\alpha - \beta\|^2$. Alice and Bob divide α and β into b blocks of size a:

$$\alpha_i = (x_{ia-a+1}, \ldots, x_{ia}), \qquad \beta_i = (y_{ia-a+1}, \ldots, y_{ia}), \qquad i = 1, \ldots b.$$

The protocol runs as follows. Alice and Bob sample b independent random vectors $U_1, \ldots, U_b$, each of them according to the distribution $\mathcal{US}(a)$. Then Alice computes b numbers

$$\langle \alpha_1, U_1 \rangle, \ldots, \langle \alpha_b, U_b \rangle,$$

and sends their approximations to Bob. More specifically, let r_i be the closest to $\langle \alpha_i, U_i \rangle$ number in $\left\{ \frac{m}{n^3} \mid m \in \mathbb{Z} \right\}$. Note that

$$|r_i - \langle \alpha_i, U_i \rangle| \leq \frac{1}{n^3}. \tag{3}$$

Alice sends $r_1, \ldots, r_b$ to Bob, each number specified by $O(\log n)$ bits. Bob computes

$$T' = (r_1 - \langle \beta_1, U_1 \rangle)^2 + \ldots + (r_b - \langle \beta_b, U_b \rangle)^2,$$

If $T' > L + \frac{5}{n}$, then Bob sends 1 to Alice. Otherwise Bob sends 0 to Alice.

Communication complexity of the protocol is $O(b \log n)$. Now we have to estimate error probability. We first show that $T' \leq L + \frac{5}{n}$ whenever $H(x,y) \leq L$ and thus the protocol does note err in this case. To this end consider the random variable

$$T = \langle \alpha_1 - \beta_1, U_1 \rangle^2 + \ldots + \langle \alpha_b - \beta_d, U_b \rangle^2.$$

Note that

$$\begin{aligned} H(x,y) = \|\alpha - \beta\|^2 &= \|\alpha_1 - \beta_1\|^2 + \ldots + \|\alpha_b - \beta_b\|^2 \\ &\geq \langle \alpha_1 - \beta_1, U_1 \rangle^2 + \ldots + \langle \alpha_b - \beta_b, U_b \rangle^2 = T. \end{aligned}$$

Let us show that $|T' - T|$ is at most $\frac{5}{n}$. Denote $P_i = r_i - \langle \beta_i, U_i \rangle$ and $Q_i = \langle \alpha_i - \beta_i, U_i \rangle$. By definition

$$T' = \sum_{i=1}^{b} P_i^2, \qquad T = \sum_{i=1}^{b} Q_i^2.$$

Thus $|T' - T| \leq \sum\limits_{i=1}^{b} |P_i^2 - Q_i^2| = \sum\limits_{i=1}^{b} |P_i - Q_i| \cdot |P_i + Q_i|$. Let us bound $|P_i - Q_i|$ and $|P_i + Q_i|$ separately. By (3) $|P_i - Q_i| \leq \frac{1}{n^3}$. By definition $|P_i + Q_i|$ is at most

$$\begin{aligned} |P_i + Q_i| &= |r_i + \langle \alpha_i, U_i \rangle - 2\langle \beta_i, U_i \rangle| \\ &\leq |r_i| + |\langle \alpha_i, U_i \rangle| + 2\,|\langle \beta_i, U_i \rangle| \\ &\leq 2\,|\langle \alpha_i, U_i \rangle| + \frac{1}{n^3} + 2\,|\langle \beta_i, U_i \rangle| \end{aligned}$$

(again we use that r_i is $\frac{1}{n^3}$-close to $\langle \alpha_i, U_i \rangle$). Coordinates of α_i and β_i are zeros and ones and there are at most n ones among them. Hence

$$|\langle \alpha, U_i \rangle| \leq \|\alpha_i\| \leq \sqrt{n}, \qquad |\langle \beta_i, U_i \rangle| \leq \|\beta_i\| \leq \sqrt{n},$$

and therefore $|P_i + Q_i| \leq 4\sqrt{n} + \frac{1}{n^3} \leq 5n$. Finally

$$|T - T'| \leq \sum_{i=1}^{b} |P_i - Q_i| \cdot |P_i + Q_i| \leq b \cdot \frac{1}{n^3} \cdot 5n \leq \frac{5}{n},$$

since $b \leq n$.

Assume that $H(x, y) = \|\alpha - \beta\|^2 \leq L$. Then

$$T' \leq T + \frac{5}{n} \leq L + \frac{5}{n}.$$

In this case the protocol always outputs 0.

Now assume that $H(x, y) = \|\alpha - \beta\|^2 \geq U$. We will show that event $T' \leq L + \frac{5}{n}$ happens with small probability. By Lemma 1 we have that

$$\mathrm{E}T = \frac{\|\alpha_1 - \beta_1\|^2}{a} + \ldots + \frac{\|\alpha_b - \beta_b\|^2}{a} = \frac{\|\alpha - \beta\|^2}{a} = \frac{H(x, y)}{a} \geq \frac{U}{a}.$$

For each $i = 1, \ldots b$ we have

$$\langle \alpha_i - \beta_i, U_i \rangle^2 \in \left[0,\, \|\alpha_i - \beta_i\|^2\right]$$

with probability 1. To finish the proof we use the Hoeffding inequality:

Proposition 3 ([6]). *If random variables $\chi_1, \ldots, \chi_m$ are independent and for each $i = 1, \ldots m$*

$$\chi_i \in [a_i, b_i]$$

with probability 1, then for every positive δ

$$\Pr\left[\chi_1 + \ldots + \chi_m \leq \mathrm{E}(\chi_1 + \ldots + \chi_m) - \delta\right] \leq \exp\left\{ -\frac{2\delta^2}{\sum\limits_{i=1}^{m} (b_i - a_i)^2} \right\}.$$

The following chain of inequalities finishes the proof.

$$\Pr\left[T' \le L + \frac{5}{n}\right] \le \Pr\left[T \le L + \frac{10}{n}\right] \tag{4}$$

$$\le \Pr\left[T \le \mathrm{E}T - ET/2\right] \tag{5}$$

$$\le \exp\left\{-\frac{(\mathrm{E}T)^2}{2(\|\alpha_1 - \beta_1\|^4 + \ldots + \|\alpha_b - \beta_b\|^4)}\right\} \tag{6}$$

$$\le \exp\left\{-\frac{\|\alpha - \beta\|^4/a^2}{2a(\|\alpha_1 - \beta_1\|^2 + \ldots + \|\alpha_b - \beta_b\|^2)}\right\} \tag{7}$$

$$\le \exp\left\{-\frac{U}{2a^3}\right\} \tag{8}$$

$$\le \exp\{-s\}. \tag{9}$$

Let us explain it step by step.

- (4) holds because $|T - T'| \le \frac{5}{n}$;
- First of all, by definition of b and since $s \ge \frac{\left(L+\frac{10}{n}\right)^3}{U^2}$ we have

  $$b^3 \ge \frac{4^3 s n^3}{U} \ge \frac{4^3 \frac{\left(L+\frac{10}{n}\right)^3}{U^2} n^3}{U} = \left(\frac{4n\left(L+\frac{10}{n}\right)}{U}\right)^3,$$

  and hence $b \ge \frac{4n\left(L+\frac{10}{n}\right)}{U}$. Now recall that $\mathrm{E}T \ge \frac{U}{a}$ and by (2) $ab \le 2n$. Therefore

  $$\frac{ET}{2} \ge \frac{U}{2a} \ge \frac{bU}{4n} \ge L + \frac{10}{n}$$

  and (5) follows.
- (6) holds because of Hoeffding inequality, applied to

  $$T = \langle \alpha_1 - \beta_1, U_1\rangle^2 + \ldots + \langle \alpha_b - \beta_b, U_b\rangle^2.$$

- (7) holds because $\|\alpha_i - \beta_i\|^2 \le a$ and hence

  $$\|\alpha_i - \beta_i\|^4 \le a\|\alpha_i - \beta_i\|^2.$$

- (8) holds because

  $$\|\alpha_1 - \beta_1\|^2 + \ldots + \|\alpha_b - \beta_b\|^2 = \|\alpha - \beta\|^2$$

  and $\|\alpha - \beta\|^2 \ge U$.
- For (9) again recall that $a \le \frac{2n}{b}$ and $b^3 \ge \frac{4^3 s n^3}{U}$.

  $$\frac{U}{2a^3} \ge \frac{U}{2\left(\frac{2n}{b}\right)^3} = \frac{b^3 U}{16 n^3} \ge s.$$

□

4 Deterministic Communication Complexity of $\mathsf{GHD}^n_{0,\,t}$

4.1 Proof of Theorem 3

Observe that

$$\mathrm{D}(\mathsf{GHD}^n_{0,\,n}) = 2 = n - \log_2 V_2\left(n, \left\lfloor \frac{n}{2} \right\rfloor\right) + O(\log n).$$

Hence we can assume that $t < n$.

Consider the protocol π witnessing $\mathrm{D}(\mathsf{GHD}^n_{0,\,t})$. Let $\pi(x, y)$ denote the leaf in protocol in which Alice and Bob come when Alice have x on input and Bob has y on input. If l is a 0–leaf of the protocol π, consider the set

$$A_l = \{x \in \{0,1\}^n \mid \pi(x, x) = l\}$$

Note that $\mathrm{diam}(A_l) \leq t - 1$. Indeed, assume that $x, y \in A_l$ and $H(x, y) \geq t$. At the same time

$$\pi(x, x) = l, \qquad \pi(y, y) = l \implies \pi(x, y) = l.$$

This contradicts the fact that l is a 0–leaf of π. Observe that

$$t - 1 = 2\left(\frac{t}{2} - \frac{1}{2}\right) \leq 2\left\lfloor \frac{t}{2} \right\rfloor, \qquad n \geq t + 1 \geq 2\left\lfloor \frac{t}{2} \right\rfloor + 1.$$

Hence by Theorem 5

$$|A_l| \leq V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right).$$

If both parties have the same $x \in \{0,1\}^n$ on input, they must come to some 0–leaf. Hence

$$\{A_l \mid l \text{ is a 0–leaf of } \pi\ \}$$

is a covering of $\{0,1\}^n$. This covering has size at most $2^{CC(\pi)}$ and each set of the covering has size at most $V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right)$. Therefore

$$2^{CC(\pi)} V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right) \geq 2^n$$

and

$$\mathrm{D}(\mathsf{GHD}^n_{0,\,t}) = CC(\pi) \geq n - \log_2 V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right).$$

Let us prove the upper bound on $\mathrm{D}(\mathsf{GHD}^n_{0,\,t})$. Let C be the covering code of radius $\left\lfloor \frac{t-1}{2} \right\rfloor$ and size at most

$$O\left(\frac{n2^n}{V_2\left(n, \left\lfloor \frac{t-1}{2} \right\rfloor\right)}\right),$$

existing by Proposition 1.

Alice computes

$$c = \arg\min_{z \in C} H(z, x)$$

and sends c to Bob. Since $c \in C$, it takes at most

$$\log_2 |C| + 1 = \log_2 O\left(\frac{n2^n}{V_2\left(n, \left\lfloor \frac{t-1}{2} \right\rfloor\right)}\right) = n - \log_2 V_2\left(n, \left\lfloor \frac{t-1}{2} \right\rfloor\right) + O(\log n)$$

bits. If $H(c, y) \leq \left\lfloor \frac{t-1}{2} \right\rfloor$, Bob sends 0 to Alice. Otherwise, Bob sends 1 to Alice.

Let us prove that the described protocol computes $\mathsf{GHD}^n_{0,\,t}$. Note that by definition of c and C we have $H(c, x) \leq \left\lfloor \frac{t-1}{2} \right\rfloor$. Hence if $x = y$, then $H(c, y) \leq \left\lfloor \frac{t-1}{2} \right\rfloor$. Assume now that $H(x, y) \geq t$. Then

$$H(x, c) + H(y, c) \geq H(x, y) \geq t > 2\left\lfloor \frac{t-1}{2} \right\rfloor$$

and hence

$$H(y, c) > 2\left\lfloor \frac{t-1}{2} \right\rfloor - H(x, c) \geq \left\lfloor \frac{t-1}{2} \right\rfloor.$$

Observe that

$$\begin{aligned} V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right) &\leq V_2\left(n, \left\lfloor \frac{t-1}{2} \right\rfloor\right) + \binom{n}{\lfloor \frac{t}{2} \rfloor} \\ &= V_2\left(n, \left\lfloor \frac{t-1}{2} \right\rfloor\right) + \binom{n}{\lfloor \frac{t}{2} \rfloor - 1} \cdot \frac{n - \lfloor \frac{t}{2} \rfloor}{\lfloor \frac{t}{2} \rfloor} \\ &\leq (1+n) V_2\left(n, \left\lfloor \frac{t-1}{2} \right\rfloor\right). \end{aligned}$$

Therefore the communication complexity of the protocol is at most

$$n - \log_2 V_2\left(n, \left\lfloor \frac{t-1}{2} \right\rfloor\right) + O(\log n) \leq n - \log_2 V_2\left(n, \left\lfloor \frac{t}{2} \right\rfloor\right) + O(\log n).$$

4.2 Application to the Number of Distinct Elements (Proof of Theorem 4)

Let F_0 denote the number of distinct elements in a given data stream of size $2n$ with tokens drawn from the universe $U = \{1, 2, \ldots, 2n\}$. We say that a deterministic p-pass streaming algorithm A with memory S for computing F_0 is *c–approximate* if A outputs a number E such that $F_0 \leq E < cF_0$. We claim that for $c < 2$ A requires $\Omega\left(\frac{n(2-c)^2}{p}\right)$ memory. Let us start with the case $p = 1$.

The first result of that kind was proved in [1]. It states that if $|E - F_0| < cF_0$, where $c = 0.1$, then A requires $\Omega(n)$ memory. A linear lower bound for memory for a larger c can be obtained by reduction to the deterministic communication

complexity of equality, as it done, for example, in [3]. Indeed, for each $\alpha < \frac{1}{2}$ there is an error-correcting code $\mathsf{ECC} : \{0,1\}^k \to \{0,1\}^n$ with relative distance at least α and $k = \Omega_\alpha(n)$. Assume that Alice has $x \in \{0,1\}^k$ and Bob has $y \in \{0,1\}^k$ their inputs. They want to decide whether $x = y$. Alice and Bob transform their inputs into 2 data streams u and v:

$$u = \langle u_1, u_2, \ldots, u_n \rangle, \qquad v = \langle v_1, v_2, \ldots, v_n \rangle,$$

where

$$u_i = n \cdot \mathsf{ECC}(x)_i + i, \qquad v_i = n \cdot \mathsf{ECC}(y)_i + i. \tag{10}$$

Alice emulates A on u. Then, using S bits, she sends a description of the current state of A to Bob and Bob emulates A on v, starting with the state he received from Alice. Finally Bob knows the output of A for the stream that is equal to the concatenation of u and v. Notice that the number of the distinct elements F_0 in this concatenation equals $n + H(\mathsf{ECC}(x), \mathsf{ECC}(y))$. If A is a $(1+\alpha)$–approximate (that is, $c = 1 + \alpha$), then Bob is able to decide whether $x = y$ or not. Indeed, if $x = y$, than $E < cF_0 = (1+\alpha)n$. If $x \neq y$, then by definition of ECC we have that $E \geq F_0 \geq n + \alpha n = (1+\alpha)n$. As deterministic communication complexity of the equality predicate on k-bit strings is k, a linear lower bound $S = \Omega(k) = \Omega(n)$ for the space complexity of A for $c < \frac{3}{2}$ follows.

In this argument we only needed a linear lower bound for 1-round communication complexity of equality predicate, which is trivial. However for arbitrary p we already need a linear lower bound for complexity of equality predicate for $2p$-round protocols. The lower bound for the space complexity we obtain by this argument becomes $\Omega(n/p)$, as in each round Alice in Bob exchange S bits.

Instead of binary error–correcting codes one can use error–correcting codes with a larger alphabet and relative distance close to 1. The same reduction provides a linear lower bound for $c < 2$. The point is that the size of the universe increases and the problem becomes harder.

Theorem 3 implies a linear lower bound on the space complexity of A for $c < 2$ in the case when the size of the universe and the size of a data stream are equal. Indeed, assume that Alice has $x \in \{0,1\}^n$ and Bob has $y \in \{0,1\}^n$, as their inputs. Assume also that they are promised that either $x = y$ or $H(x,y) \geq t = \lceil n(c-1) \rceil$. Again, Alice and Bob transform their input into data streams u and v but with (10) replaced by

$$u_i = nx_i + i, \qquad v_i = ny_i + i.$$

The expression for F_0 becomes $F_0 = n + H(x,y)$. Thus Alice and Bob can solve $\mathsf{GHD}^n_{0,\,t}$ using $2pS$ bits of communication. Indeed, if $x = y$, then $E < cF_0 = cn \leq n + t$ since by definition $t \geq (c-1)n$. If $H(x,y) \geq t$, then $E \geq F_0 \geq n + t$.

We conclude that by theorem 3 that pS must be at least

$$pS = \Omega\left(n - \log_2 V_2\left(n, \frac{t}{2}\right) + \log n\right) = \Omega\left(\left(\frac{1}{2} - \frac{t}{2n}\right)^2 n\right) = \Omega(n(2-c)^2).$$

References

[1] Alon, N., Matias, Y., and Szegedy, M. The space complexity of approximating the frequency moments. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (1996), ACM, pp. 20–29.

[2] Brody, J., Chakrabarti, A., Regev, O., Vidick, T., and De Wolf, R. Better gap-hamming lower bounds via better round elimination. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 2010, pp. 476–489.

[3] Chakrabarti, A. Data stream algorithms. *Computer Science 49*, 149.

[4] Chakrabarti, A., and Regev, O. An optimal lower bound on the communication complexity of gap-hamming-distance. *SIAM Journal on Computing 41*, 5 (2012), 1299–1317.

[5] Cohen, G., Honkala, I., Litsyn, S., and Lobstein, A. *Covering codes*, vol. 54. Elsevier, 1997.

[6] Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association 58*, 301 (1963), 13–30.

[7] Indyk, P., and Woodruff, D. Tight lower bounds for the distinct elements problem. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on* (2003), IEEE, pp. 283–288.

[8] Marsaglia, G. Choosing a point from the surface of a sphere. *The Annals of Mathematical Statistics 43*, 2 (1972), 645–646.

[9] Sherstov, A. A. The communication complexity of gap hamming distance. *Theory of Computing 8*, 1 (2012), 197–208.

[10] Sudan, M. *Algorithmic Introduction to Coding Theory: Lecture Notes*. 2001.

[11] Vidick, T. A concentration inequality for the overlap of a vector on a large set, with application to the communication complexity of the gap-hamming-distance problem. *Chicago Journal of Theoretical Computer Science 1* (2012).